On Xenakis's Games of Musical Strategy

Stefano Kalonaris
RIKEN Center for Advanced Intelligence Project (AIP)
stefano.kalonaris@riken.jp
Orcid: 0000-0002-7372-323X
DOI: 10.46926/musmat.2022v6n2.56-71

Abstract: *In this article, Xenakis's trilogy of pieces based on zero-sum games is investigated, revealing axiomatic inconsistencies between game-theoretical models and their musical translations implemented according to the aesthetic preferences of the composer. The problematic elements regard 1) the formal definition of the games and 2) the limits of objective utility functions in the music domain. After introducing these games of musical strategy and the few existing practical implementations found in the literature, a detailed comparison between the mathematical models and Xenakis's renditions is presented to highlight their divergence. Then, the feasibility from a musical perspective of the rational decision-making required by the models is explored through a computational simulation of game dynamics using Reinforcement Learning. Lastly, the article concludes by contextualising the findings to the increasingly ubiquitous role of the machine in the creative processes of musical composition and generation. In doing so, Xenakis's original works offer themselves as a springboard for (re-)imagining and developing novel approaches integrating human-machine decision processes with musical design and interaction.*

Keywords*: Iannis Xenakis, Game Theory, Reinforcement Learning.*

I Introduction

Duel (1959), *Stratégie* (1962) and *Linaia-Agon* (1972) form a trilogy of pieces that Xenakis wrote using identical underlying principles. Such principles imply *external* rather than *internal* conflict, with the latter relating to intrinsic factors (*i.e.*, the dialogical relation between the sound rendition and the symbolic schema) and the former involving extrinsic factors. Xenakis calls the music originating from these *heteronomous* (external conflict) and *autonomous* (internal conflict). These three works can be classified as *game pieces* [7, 9], which would draw inspiration from a notion of "game" linked to ludic activities, social theory, anthropology and game design. For example, John Zorn's *Cobra* [4] focuses on emergent social dynamics, Mauricio Kagel's *Match* [11] is inspired by tennis, and Mathius Shadow-Sky's *Ludus Musicae Temporarium* [22] is based on the work of Huizinga [10] and Caillois [6].

Xenakis's trilogy, however, differs in that it is stricter in its application of formal and rational theories on decision-making to the music domain. More specifically, it is rooted in the applied

Received: October 27th, 2022
Approved: December 21st, 2022

mathematics field known as *game theory*, which models scenarios of conflict or cooperation between agents. While game theory is popular in economics, social network theory, and computer science (with applications ranging from artificial intelligence to network systems), game-theoretical musical pieces remain relatively under-explored. Besides some sporadic experiments with Bayesian games of imperfect information in the context of networked music performance [14, 13, 12], game-theoretical formalisms are not a popular paradigm for composition or sound design, because of the difficulty of mapping objective utility functions in the music domain. That is, abiding by the principles of optimisation, rational decision-making and mathematical solutions might be, at times, undesirable if in conflict with musical and aesthetic goals. Imaginary rewards (*e.g.*, "points" won or lost in a game) may be less motivating than sonic, experiential counterparts (*e.g.*, the perceived quality of the music or the musical interactions originating from the musical game piece).

Xenakis's game pieces are not devoid of these issues, and it is important to consider them in the context of the (then) newly re-defined role of the composer: caught in between "inventing schemes (previously forms) and exploring the limits of these schemes", and "effecting the scientific synthesis of the new methods of construction and of sound emission" [32, p.133]. While uncompromisingly and single-mindedly striving for novelty and originality, Xenakis retained an executive role rooted in the idea of the composer as the sole owner of the work. His indisputable aesthetic compass, for example, is asserted when stating that "the winner has won simply because he has better followed the rules imposed by the composer, who, by consequence, claims all responsibility for the 'beauty' or 'ugliness' of the music" [31].

However, game theory has an aesthetic of its own. In fact, this might be the true *raison d'etre* of these works, as Xenakis confesses that he had "[. . .] been interested in social questions, in the relationship between people and the aesthetic aspect of all that" [28, p.49]. One might say that Xenakis's games of musical strategy find their most defining characteristic to be this tension between internal (the composer's quest for aesthetic integrity) and external conflict (opposing interests which must be allowed to emerge).

When revisiting Xenakis's games of musical strategy, it is also crucial to acknowledge the changes brought forward by the increasingly common application of AI techniques and methods in the music domain and how these retroactively affect one's understanding and appreciation of Xenakis's game pieces. According to Xenakis, the role of the machine is perceived either through negative or positive bias or as an explorative process which, however, is ultimately not sufficient *per se* as a means to artistic value [32]. In agreement with this viewpoint, this article is, nevertheless, focused on analyses and implementations of Xenakis's trilogy comprising a computer-assisted factor.

I.i Background

Despite their simplicity and potential for myriad variations and implementations, Xenakis's game pieces have not enjoyed as much attention as the remainder of his body of work. Liuni & Morelli [18] rendered *Duel* as a live installation in which members of the public could take the role of the conductors. Their movements were analysed by computer vision algorithms to drive the score, which was, in turn, visualised on a large screen. This rendition concentrated on the interactive participation of the human players, and the musical events prescribed in the original score were realised using audio samples instead of real orchestras. Regarding *Linaia-Agon*, there exists a documentary DVD [24] with original radio broadcasts and newly recorded live and studio performances of the piece using a computational interface.

As for the computational analysis of Xenakis's trilogy, Sluchin & Malt [25] simulated *Duel* based on four different methods for choosing tactics (see Section II.iii). Beyond the numerical simulation of game dynamics, that work, and a follow-up study [26] which included a computational interface to switch between different strategies, does not point to an audio rendition of the piece. *Linaia-Agon*, finally, was analysed by DeLio [8], Sluchin [23], and Beguš [3].

The next section precedes the formal analysis of Xenakis's trilogy by introducing the reader to some fundamental notions in game theory.

II Fundamental Notions

In the context of game theory, "a game is a description of strategic interaction that includes the constraints on the actions that the players can take and the players' interests, but does not specify the actions that the players do take" [21, p.2]. The basic elements of a game are:

- a finite set of **players** N
- for each player $i \in N$ a non-empty set A_i representing the set of **actions** available to player i
- for each player $i \in N$ a **preference relation** $\succsim_i$ on $A = \times_{j \in N} A_j$ (the set of outcomes by A)

Morgenstern & von Neumann [19] are credited as the initiators of modern game theory, which is normally divided into two main branches: **non-cooperative** and **cooperative** game theory. The former considers each player's individual actions as primitives, whereas the latter sees joint actions as primitives, and assumes that binding agreements can be made by players and within groups of players. For the sake of this article, only non-cooperative games are discussed.

Another fundamental characterisation of games that will be crucial for Xenakis's game pieces is based on the **utility function** which maps rewards to actions $u_i : A \to \mathbb{R}$, so that $u_i(a) \geq u_i(b)$ whenever $a \succsim_i b$. These "rewards" are hereinafter referred to as **payoffs**. To this end, one can distinguish between **constant-sum** and **variable-sum** games. In the former, the payoffs for each possible combination of actions sum up to the same constant C. A particular case of constant-sum games is **zero-sum** games. This means, simply, that for every combination of players' actions, one player's gains are the other's losses, and therefore payoffs sum up to 0. In variable-sum games, on the other hand, the rewards are neither symmetrical (opposites) nor sum up to a constant.

A **strategy** is a decision algorithm which considers the options available under a given scenario. In other words, a strategy is a complete contingent plan that defines the action a player will take in all states of the game. A **strategy profile** is a set of strategies for all players that fully specify all actions in a game.

Insofar as the discussions of this article are concerned, the most important aspects of a game entail notions of *time*, *information*, and *equilibrium*.

II.i Time

This dimension determines whether games are *simultaneous* or *sequential*, meaning whether players take actions synchronously and independently from each other, or in turns, respectively. Simultaneous and sequential games are normally notated differently. For the former, one uses the **normal form** (or **strategic form**), while for the latter the **extensive form** (or **game tree**). For simplicity, two players are considered, hereinafter x and y. Examples of two games and their graphical representations can be seen in Figure 1. In zero-sum games, such as the game shown on the left (known as *Rock, Paper, Scissors*), for each cell in the matrix, payoffs are expressed as a single signed integer. Since this is a zero-sum game, the payoffs are symmetrical (opposites); for example, -1 means $(-1, 1)$ where the first value in the tuple would be assigned to x and the second to y. Instead, in the sequential game depicted on the right, the payoffs are not symmetrical and are explicitly expressed as tuples of values.

		y: Rock	y: Paper	y: Scissors
x	Rock	0	-1	1
	Paper	1	0	-1
	Scissors	-1	1	0

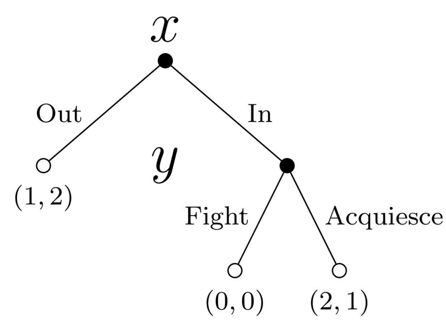

Figure 1: *Simultaneous (left) and sequential (right) games, and their notation: normal and extensive form, respectively.*

II.ii Information

In this dimension, one must distinguish between **perfect** *vs.* **imperfect** and **complete** *vs.* **incomplete** information. The former describes whether or not players have knowledge of each others' actions and history. The latter is instead concerned with common knowledge of each player's utility functions, payoffs, strategies and "types". The two are not mutually exclusive: for example, a game could have perfect and incomplete information, and so forth. These distinctions are not trivial and fundamentally affect the decision-making process.

It is possible to convert simultaneous games from normal to extensive form, and vice versa (**induced normal form**). To illustrate this procedure, a zero-sum game known as *Matching Pennies* is considered. To convert it to extended form using the game matrix, one introduces an **information set**, indicated as a dotted ellipse enveloping decision leaf nodes. The information set shown in the game tree on the right in Figure 2 indicates that player y, while moving after player x, has no knowledge of what action the opponent chose and, therefore, whether she finds herself under the right or left leaf node. This uncertainty would define such a game as an extensive form game of incomplete information.

		y: heads	y: tails
x	heads	1	-1
	tails	-1	1

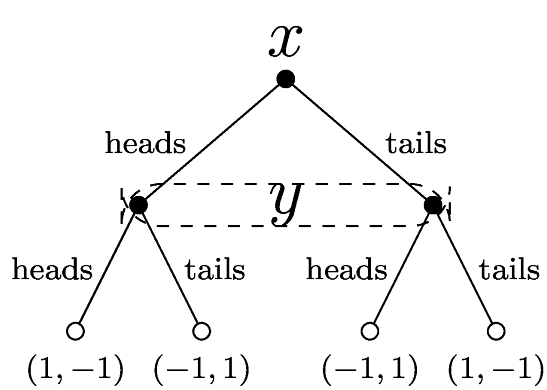

Figure 2: *A popular zero-sum game known as* Matching Pennies*, shown in matrix form (left), and in an equivalent extensive form (right).*

However, this article is only concerned with games with perfect and complete information, either simultaneous or sequential. From now on, the former will be referred to simply as **strategic games**, whereas extensive-form games with complete information will be referred to as **extensive**

games. For the latter, two more elements are needed for a formal definition:

- a set H of sequences where each member is a **history** and each component of a history is an action. A history $(a^k)_{k=1,\ldots,K} \in H$ is **terminal** if it is infinite or if there is no a^{K+1} such that $(a^k)_{k=1,\ldots,K+1} \in H$. The set of terminal histories is indicated as Z
- a **player function** P that assigns to each non-terminal history (each member of $H \setminus Z$) a member of N

II.iii Equilibrium

There are several approaches to solving games. Solutions are optimal combinations of strategies that ensure the best outcome for a given or all players.

For strategic games, different game-theoretic solution concepts include **maximax** (maximise one's payoff), **maximin** (maximise one's minimum payoff, a.k.a. choosing the best of the worst possible outcomes), and **minimax** (minimise one's maximum loss). In a two-player zero-sum game, when the matrix has a **saddle point**, meaning a given action pair yields the best outcome for both players (*i.e.*, neither could do any better), the maximin and minimax strategies produce the same result. Thus, one can define **Nash equilibrium** for a strategic game $\langle N, (A_i), (\succsim_i)\rangle$ as the solution where no player has an incentive to change strategy given that no one else does, and express it formally as a profile $a^* \in A$ for every player $i \in N$ so that $(a^*_{-i}, a^*_i) \succsim_i (a^*_{-i}, a_i)$ for all $a_i \in A_i$.

When there is no saddle point, one can use the notion of equilibrium defined in terms of **mixed strategies** [19]. This involves randomising one's action selection with weighted probabilities, which ensure statistically optimal outcomes. In other words, each player makes the other indifferent between choosing one action or another, so neither player has an incentive to try another strategy. As a practical example, one can use the previously seen game of *Matching Pennies*. One assumes that player x plays heads with probability p and tails with probability $1-p$. Similarly, player y plays heads with probability q and tails with probability $1-q$. According to the payoff values in Figure 2, x's rewards will be $1 \cdot q - 1 \cdot (1-q)$ for playing heads and $-1 \cdot q + 1 \cdot (1-q)$ for playing tails. In equilibrium, player x is willing to randomise only when she is indifferent between heads and tails. Thus, the two equations must be equal, yielding $q = \frac{1}{2}$. Following identical reasoning for player y, one obtains $p = \frac{1}{2}$. Therefore, both players will play heads or tails with a probability of $\frac{1}{2}$. This is a simple if trivial example, but mixed strategies, if allowed, can always guarantee an equilibrium.

For an extensive game $\langle N, H, P, (\succsim_i)\rangle$, and having defined $O(s)$ as the outcome of a strategy profile $s = (s_i)_{i \in N}$, the Nash equilibrium will be the strategy profile s^* for every player $i \in N$ such that $O(s^*_{-i}, s^*_i) \succsim_i O(s^*_{-i}, s_i)$ for every strategy s_i of player i. Mixed strategies can work analogously to what is seen in strategic games, whereby a Nash equilibrium in mixed strategies for extensive games can be expressed as a profile σ^* of mixed strategies so that $O(\sigma^*_{-i}, \sigma^*_i) \succsim_i O(\sigma^*_{-i}, \sigma_i)$ for every mixed strategy σ_i of player i.

However, because of the sequential nature of extensive games, the notion of **subgame perfect equilibrium** is introduced. This is a refinement of the Nash equilibrium defined above which accounts for history-dependent best responses so that, for every non-terminal history $h \in H \setminus Z$ for which the player function is $P(h) = i$, $O(s^*_{-i}|_h, s^*_i|_h) \succsim_i O(s^*_{-i}|_h, s_i|_h)$ for every strategy s_i of player i, for the subgame $G(h)$. Normally, subgame perfect equilibrium is obtained using **backwards induction**: starting from the terminal history, one finds the best response strategy profiles or the Nash equilibria in the subgame, assigns these strategy profiles and the associated payoffs to the subgame, and moves successively towards the beginning of the game.

The next section looks at how Xenakis leveraged game-theoretical concepts to design musical game pieces.

III Xenakis's Game Pieces

Xenakis's trilogy is based on zero-sum games with complete and perfect information: each player, when making a decision, has knowledge of all the events that have previously occurred (*i.e.*, actions taken by the opponent are observable). Xenakis expresses all three games as strategic games using the normal form, and he provides mixed strategy calculations since these game matrices have no saddle points.

The details for each zero-sum game piece in Xenakis's trilogy follow.

III.i *Duel*

Duel is described in Chapter IV of *Formalized Music* [32], where the reader is referred for finer details. This game piece sees two conductors and their respective orchestras (hereinafter x and y, in keeping with the convention adopted thus far) competing against each other via means of juxtaposing musical events (or *tactics*, in Xenakis's choice of terms). Said tactics are chosen based on combinations of pairwise actions (*i.e.*, the payoff matrix) associated with an aesthetic outcome value stipulated by the composer. Details on the musical instructions for these events are omitted here for the sake of brevity. The payoff matrix undergoes several transformations to ensure a fair game, and its final form is shown in Figure 3.

x \ y	I	II	III	IV	V	VI	
I	-1	+1	+3	-1	+1	-1	$\frac{14}{56}$
II	+1	-1	-1	-1	+1	-1	$\frac{6}{56}$
III	+3	-1	-3	+5	+1	-3	$\frac{6}{56}$
IV	-1	+3	+3	-1	-1	-1	$\frac{6}{56}$
V	+1	-1	+1	+1	-1	-1	$\frac{8}{56}$
VI	-1	-1	-3	-1	-1	+3	$\frac{16}{56}$
	$\frac{19}{56}$	$\frac{7}{56}$	$\frac{6}{56}$	$\frac{1}{56}$	$\frac{7}{56}$	$\frac{16}{56}$	

Figure 3: Duel*'s payoff matrix and associated weights of the mixed strategies.*

The mixed strategy Nash equilibrium for *Duel* can be calculated by determining the probability corresponding to each strategy so that x is indifferent to the actions of y, and vice versa. Each conductor might select the next event based purely on these probabilistic weights. For example, x will pick event I with a probability of $\frac{14}{56}$, and so forth.

III.ii *Stratégie*

Like in *Duel*, there are 6 possible tactics to choose from, but their musical content differs in *Stratégie*. Conceptually an extension of *Duel*, *Stratégie* augments the scope of the score by allowing combinations of 2 or 3 tactics for each conductor. These yield 19 possible combinations for each conductor to choose from and 361 possible pairs of such composite events. The resulting matrix is omitted here. In addition to this matrix, *Stratégie* also provides two additional 3×3 matrices that, by aggregating two-by-two and three-by-three compatible combinations of fundamental tactics, reduce the decisional complexity for the conductors.

x \ y	▼ • H ∴ # III	•▼ H▼ ∴▼ #▼ III▼ H• ∴• #• III•	H•▼ ∴•▼ #•▼ III•▼	
▼ • H ∴ # III	3	-1	0	$\frac{1}{8}$
•▼ H▼ ∴▼ #▼ III▼ H• ∴• #• III•	-5	-1	4	$\frac{3}{8}$
H•▼ ∴•▼ #•▼ III•▼	3	1	-3	$\frac{4}{8}$
	$\frac{1}{6}$	$\frac{3}{6}$	$\frac{2}{6}$	

Figure 4: *One of* Stratégie*'s* 3×3 *payoff matrices and associated weights of the mixed strategies.*

III.iii *Linaia-Agon*

This piece is inspired by mythological tales of a human musician, Linos, challenging the god Apollo. *Linaia-Agon* is a three-part work, of which the second is structured around zero-sum games of strategy, not unlike those seen so far. This part comprises two games: *Choice of Combats* and *Combats*. Figure 5 shows the former on the left and one of the combats on the right, with Linos (played by a trombone) indicated as x and Apollo (played by a tuba) marked as y, for the sake of consistency with earlier seen game matrices. In *Choice of Combats*, the available tactics are indicated as α, β, γ and correspond to specific notes (different for trombone and tuba). The sequence of their occurrence in the game then goes to determine the following part (*i.e.*, *Combats*), where each of these three Greek letters corresponds to a new game matrix, each with different payoffs (and resulting mixed strategy Nash equilibrium) and instrumentation (only for Apollo). In these three games, there are four tactics that one can choose.

Xenakis provides instructions and guidelines for playing his games of musical strategy. If carefully examined, however, these explanations, along with additional evidence gathered from interviews containing direct references or notions relating to the trilogy under scrutiny, might

		y: α	β	γ	
x	α	-3	-8	7	$\frac{4}{15}$
	β	2	2	-3	$\frac{10}{15}$
	γ	-8	12	2	$\frac{1}{15}$
		$\frac{6}{15}$	$\frac{3}{15}$	$\frac{6}{15}$	

		y: ≁	∵	/	∅	
x	≁	1	-2	-2	3	$\frac{92}{335}$
	∵	1	-1	0	-2	$\frac{78}{335}$
	/	0	4	-2	-2	$\frac{89}{335}$
	∅	-2	-1	5	1	$\frac{76}{335}$
		$\frac{161}{335}$	$\frac{61}{335}$	$\frac{72}{335}$	$\frac{41}{335}$	

Figure 5: *On the left:* Linaia-agon*'s payoff matrix for the* Choice of Combats *game. On the right:* Linaia-agon*'s β game matrix, one of the three* Combats*. Tactics symbols and corresponding musical content are as follows:* ≁ *indicates frequency and amplitude modulation,* ∵ *stands for staccato articulation,* / *means glissandi, and* ∅ *was originally left null in Xenakis's score but is here introduced to denote silence.*

reveal contradictions and incongruencies. The points of contention are discussed in the next section.

IV Insight

Although discrepancies between the theoretical axioms and the score implementations have been flagged [2] for other works of Xenakis, similar discussions concerning his game pieces have not, to the author's knowledge, been had as yet. This section aims to provide deeper insight and a critical understanding of Xenakis's games of musical strategy.

IV.i Game Formalism

Perhaps the most important issue in Xenakis's trilogy concerns the formal definition of the game model. In particular, there is little clarity with respect to the time dimension (as defined in Section II.i). In the original description of *Duel* and *Stratégie*, Xenakis states that the payoffs refer to "couples of simultaneous events" [32, p.114] and again that "pairs of tactics are performed simultaneously" [32, p.126]. However, in Figure IV-4 [32, p.126], reproduced here in Figure 6, one can see that the tactics in a pair are instead asynchronous: one conductor starts by choosing event i and the other conductor responds with event j. This is further corroborated when, after having decided who is x and who is y by means of a coin toss, "deciding who starts the game is determined by a second toss" [32, p.126].

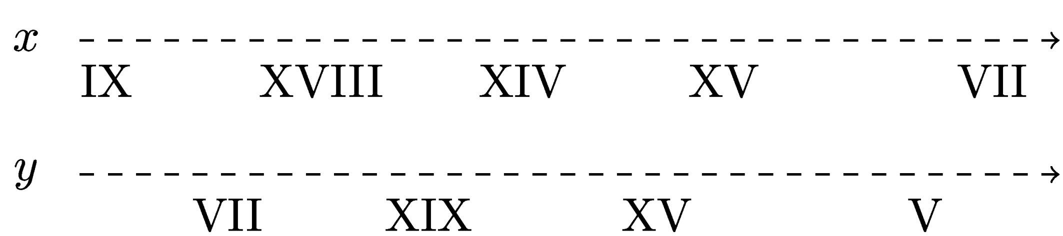

Figure 6: *A reproduction of Xenakis's original Figure IV-4 in* Formalized Music*. Note: in this version, only the temporal sequence of tactics is displayed, whereas corresponding payoffs are omitted.*

Moreover, when discussing how to award points, Xenakis suggests "to have one or two referees counting the points in two columns, one for conductor X and one for conductor Y, both in positive numbers". As already explained, however, this is a zero-sum game meaning that one's gains are the other's losses. Therefore, this instruction is somewhat confusing. Another option for awarding points is to use "an automatic system that consists of an individual board for each conductor". In this case, Xenakis explains that, for example, "if conductor X chooses tactic XV against Y's IV, he presses the button at the intersection of row XV and column IV". The use of the word "against" suggests that one conductor selects a tactic conditioned on the tactic of the opponent. So far, all evidence points to sequential games in extensive form. The choice of representation is not a reason for debate, since it is possible to convert from normal to extensive forms, as seen in Section II.i. However, univocal clarity about the time dimension of Xenakis's game pieces becomes problematic because in sequential games players who move later in the game can condition their choices on observed moves made earlier in the game. Conversely, in simultaneous games, players must all choose their own strategies without knowing what strategies are chosen by other players. These two types of games are solved differently, based on this information dependency, as seen in Section II.iii.

Xenakis's games are **repeated**, and at each repetition (hereinafter, **stage game**) both conductors are presented again with the entire game matrix, thus, effectively, starting anew. Repeated games are normally described as extensive form games where each stage game is modelled on a normal form. That is, each stage game is considered either as a strategic game or as an extensive form game with imperfect information derived from the normal form. As Xenakis's games have been shown to be sequential, it can be concluded that they should be classified as finitely repeated sequential games with perfect and complete information, or, more succinctly, **finitely repeated extensive games**. To summarise:

- The strategic relationship between players is expressed by a normal form (the payoff matrix)
- Players play an extensive form game (a sequential game) derived from the normal form
- Payoffs are handed out after every stage game
- Every stage game is the same in every stage
- One can find the minimax point using the normal form representation of the stage game and define the subgame perfect equilibrium in the corresponding repeated game

Having formally defined the game model for Xenakis's trilogy, the next step is to look at the decision-making processes that might provide viable solutions.

IV.ii Decision-making

Xenakis discusses decision-making strategies explicitly when referring to *Stratégie*, but these equally apply to *Duel*. He provides the following options:

1. arbitrary choice
2. *a priori* agreement on a sequence of action pairs
3. "drawing from an urn containing balls [. . .] in different proportions"
4. eliminating one conductor (the remaining one directs both orchestras)
5. conditioning on "the winnings or losses contained in the game matrix"

Option (3) is effectively equivalent to abiding by the mixed strategy Nash equilibrium. Arguably, probabilistic pooling is a difficult task for the average folk: randomising using weighted probabilities would require some aid, which is what Xenakis suggests (*i.e.*, urns). The composer further recommends that this be done offline, before the performance, and adequately rehearsed. However, this option, along with (1), (2) and (4), is disregarded. This is somewhat strange given that a good part of the chapter dedicated to *Duel* and *Stratègie* is spent on optimisations based on randomisation strategies. Except for the last option, all these decision-making methods are deemed unsuitable and termed *degenerate*, lacking "any conditioning for conflict, and therefore without any new compositional argument" [32, p.113]. One is thus left with subgame perfect equilibrium either via a) backward induction or b) choosing the best response at each stage game and disregarding the history up to that point. Whether to use one or the other is largely dependent on how the game is going to end. Xenakis provides three options to limit the game: based on a fixed number of stage games (histories), on a fixed cumulative payoff value or on a fixed time length.

To use backward induction to calculate the subgame perfect equilibrium, it is necessary to know the terminal history Z. There could be two ways to determine the number of histories. In one case this would be decided on the spot, just before starting the game. This would put conductors under considerable strain as they would have to be apt in rapid backward induction calculations. Otherwise, the number of histories could be determined before the start of the game, leaving ample time to calculate the subgame perfect equilibrium using backward induction *a priori*. Then, the conductors would simply follow the sequence of action pairs in H during the performance. This case would reduce the decisional power of the conductors to solely choosing a time for the next action, and it would arguably be at least as dismissive of "conditioning for conflict" as the other degenerate options listed earlier.

If the number of stage games is unknown (*i.e.*, the game ends upon reaching a stipulated cumulative payoff or time value) and conductors are advised not to use mixed strategies (see above), then they must simply select the action that yields the best outcome in each stage game. This is akin to a memoryless system since each conductor only considers the current state to deliberate. Besides offering an impoverished notion of decision-making agency (arguably not particularly skilful in dealing with expectation or insightful in the opponent's ways), purely reacting at each decision node without knowledge of the past highlights another problematic aspect of Xenakis's game design. That is, there might be times when having to choose between equivalent payoffs for different action pairs in the game matrix, as in the case of *Duel* (see Figure 3) and *Linaia-Agon*'s *Choice of Combat* (see Figure 5). How would then a conductor break these ties? Random selection must be excluded since it is *degenerate* according to Xenakis (see above). What about personal preference regarding the aesthetic value of candidate actions or action pairs? This would seem reasonable, although in conflict with the aesthetic top-down control advocated when one is told that "I am the judge - the one who determines which solution is more interesting" [28, p.108].

Finally, it must be noted that a subgame perfect equilibrium requires that, regardless of what players observe, they will continue to maintain the original assumptions that the opponent is 1) rational 2) knows the game or perceives it identically to how it has been specified and 3) does not make mistakes. If the sole aim was to keep with a less lenient approach that foregrounds the mathematical foundations of the game piece over an ill-defined utility function in the music domain, then computational conductors could be employed. To investigate this option, the next section offers a simple simulation using computational conductors that can learn optimal strategies.

V Experiment: *Duel*

Duel is considered as a case study. The experiment presented here is not a mere replication of that conducted by Sluchin & Malt [25] because it involves computational decision-making agents that are able to learn how to solve *Duel* via either minimax or mixed strategies. The game is treated as a simultaneous repeated game, allowing the conductors to continue playing it as many times as necessary to learn the best response. A conductor is modelled using Reinforcement Learning (RL), one of the three main paradigms used in machine learning (the others being: supervised and unsupervised learning) and the most widely applied to gaming problems, in general. RL involves a loop whereby agents in an environment take actions and receive feedback from some reward function, thus incrementally learning to optimise their cumulative rewards. RL is normally modelled as a Markov Decision Process (MDP): at each discrete time step t, the agent receives the current state s_t and a reward r_t, chooses an action a_t from the set of available actions A, which modifies the environment to a new state s_{t+1} and yields a new reward r_{t+1} based on the transition (s_t, a_t, s_{t+1}). Through this iterative process, the learning agent aims to learn a policy: $\pi : A \times S \to [0, 1], \pi(a, s) = \Pr(a_t = a \mid s_t = s)$ that maximises the expected cumulative reward. Arguably, RL is very suitable for modelling game-theoretical tasks [20].

V.i Q-Learning

Q-learning [30] is a model-free RL algorithm, meaning that it does not use the transition probability distribution associated with an MDP, but it is rather a trial-and-error approach. Q-learning was chosen for the experiment presented in this section because it is applicable exclusively to discrete action and state spaces, which is the case of Xenakis's game pieces. Q stands for the quality of a state-action combination, and it is expressed as $Q : S \times A \to \mathbb{R}$ using a value iteration update equation that accounts for the weighted average of the old value and the new information, as well as other parameters such as a learning rate α and a discount factor γ. Simply put, the Q-learning algorithm works as follows: initialise a Q-table ($n \times m$ with n = number of actions and m = number of states), choose an action, measure the reward, update the Q-table, repeat.

V.ii Tournament

For this experiment, three types of conductors are used: one that learns the mixed strategy Nash equilibrium using Q-learning (hereinafter, *NashQ*), another that learns the minimax strategy using Q-learning (from now on referred to as *MinimaxQ*), and yet another which does not learn but simply selects tactics at random (hence, called *Random*). The simulation was implemented using the Python[1] programming language with few additional dependencies. These included the *nashpy* library[2] which offers different algorithms for the calculation of the Nash Equilibria, such as *vertex enumeration* [29] or the Lemke & Howson [17] method. Learning conductors were instantiated with a default value of $\gamma = 0.99$. In this case, a tournament comprises all possible combinations of conductors, with repetition (*e.g.*, $[A, A], [B, B]$, etc.) but with order invariance (*e.g.*, $[A, B] == [B, A]$). One instance of a repeated game between any given pair will hereafter be referred to as an *epoch*, whereas one repetition of the game in an epoch will be called an *episode*.

Figure 7 shows one epoch in a tournament, between a NashQ and a MinimaxQ, while Figure 8 shows the duel between two NashQ. Figure 9, instead, shows a duel involving a Random conductor, as a comparison baseline. From these plots, it is possible to note that, after an initial

[1] https://python.org
[2] https://pypi.org/project/nashpy/

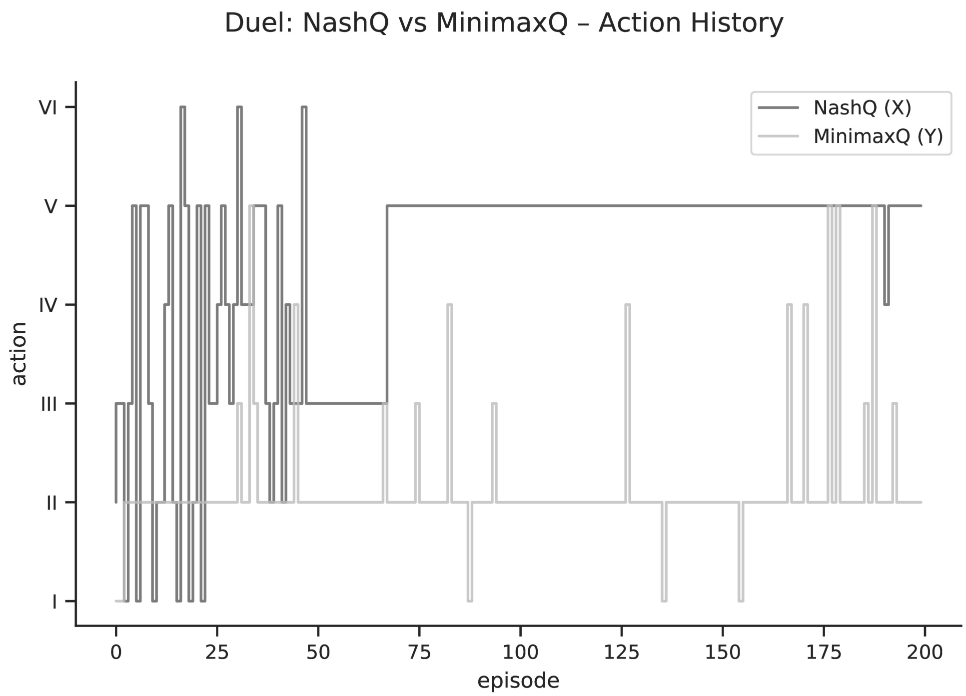

Figure 7: *Duel between a NashQ conductor and a MinimaxQ opponent, playing a repeated game of 200 episodes.*

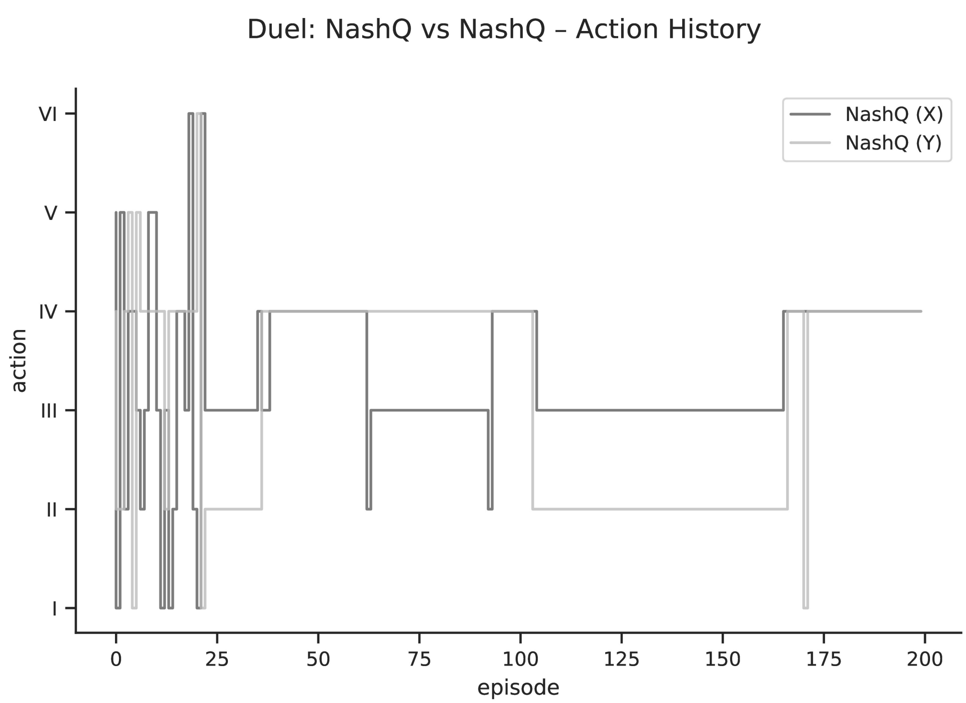

Figure 8: *Duel between two NashQ conductors, playing a repeated game of 200 episodes.*

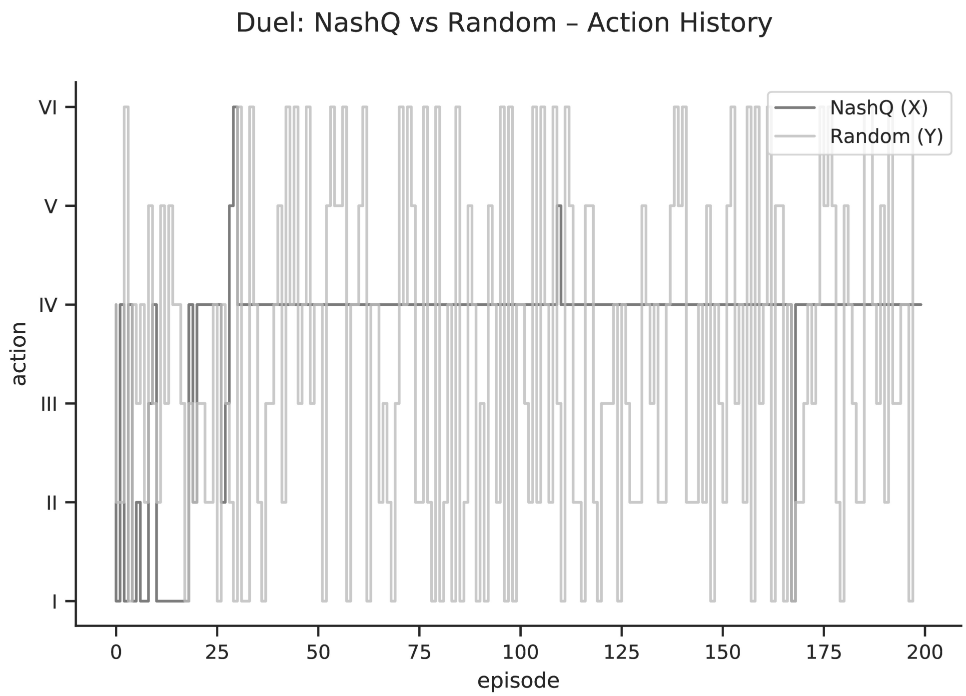

Figure 9: *Duel between a NashQ conductors and a Random conductor, playing a repeated game of 200 episodes.*

phase of learning, the conductors converge towards stable strategies, meaning that they end up playing the same tactic for many consecutive episodes.

In a real-performance scenario, how would these learnt behaviours and policies potentially affect the resulting music? Although there is general consensus on the basic principles of design in visual arts [1], the same cannot be said for musical design and composition. However, the notion of *contrast* (or variety) seems to be a constant among all the arts, including music. Using contrast as an evaluation metric, would imply that learned, informed conductors might be liable for considerably lower musical contrast (boring music, some might say?), whilst being numerically optimal.

VI Reflection & Future Work

It is hopefully clear by now how the choice of an applied mathematics framework gives rise to some interesting disjunctures. The axioms of game theory, which presuppose a well-defined utility function, can be challenging to uphold when artistic, creative, and aesthetic goals compete with theoretical ones. One must then consider that the intrinsic value of a game-based framework resides in the possibility to subvert its rational/theoretical axioms in favour of more palatable, musically pleasing, or desirable outcomes. Doing so allows x or y to make sub-optimal or biased choices in terms of payoff that are instead optimal from an artistic (subjective) viewpoint. Given the current involvement of computers in creative and musical tasks, it is important to contextualise the lessons learned so far. Solving Xenakis's game matrices is a trivial task. Simulating *naturalistic decision-making* [16] that might communicate a sense of negotiation between aesthetic or musical concerns and the mathematical imperatives of the game, on the other hand, is not. To this end, when designing computational versions of Xenakis's zero-sum games of strategy, one could deliberately introduce some sub-optimality [15], even opening up to simple options such as finite state machines. In this case, decision policies could stochastically include not only *mixed strategies* and *minimax*, but also subversive states based on the simulation of some aesthetic preference, for example, or some intrinsic attitude trait (*e.g.*, impulsive, greedy, rational, etc.). Employing simple approaches such as automata might sound naïve in an age of large language models with billions of parameters [5] and an in-depth discussion of the best approaches to simulate a process as complex as human decision-making is beyond the scope of this article (and possibly beyond the scope of machine learning techniques to date). Thus, these speculations are merely included as a springboard for further exploration of Xenakis's game pieces. For example, paradigms such as interactive evolutionary computation [27], whereby the reward function (the fitness function, in this case) is updated according to the feedback given by a human meta-conductor, could be leveraged. Furthermore, and particularly given the recent shift toward increased online presence due to the COVID-19 pandemic, telematic, remote implementations of Xenakis's game pieces could also be envisaged. These are just a few re-imagined incarnations of Xenakis's trilogy among the endless possibilities available.

VII Conclusion

This article considered the games of musical strategy composed by Xenakis between 1959 and 1972 and discussed philosophical, musical, and mathematical properties of these works. In identifying a problematic disparity between the game-theoretical models and their rendition in Xenakis's game-based works, the relationship between rational decision-making and the notion of payoff or reward in a musical context foregrounded further incongruencies. It was posited that a payoff matrix alone is not a sufficient incentive or motivation to resolve musical conflict and that, if one wishes to use applied mathematics for musical interaction in a strict fashion, disappointments regarding both numerical solutions and musical texture are likely to emerge. These claims were substantiated by an in-depth probe at a theoretical level, accompanied by supporting evidence offered in the form of direct references in Xenakis's game description and instructions. More evidence was gathered via a simulation of *Duel*, where computational learning agents took on the role of conductors and battled in an extensive tournament. Said agents were modelled using Reinforcement Learning to optimise their responses over time, based on different solutions (namely mixed strategies and minimax). An additional agent that responded arbitrarily (without learning) was also used as a comparison baseline. In summary, this article attempted to show that the experience of playing a game of musical strategy is a complex phenomenon where disparate factors converge, transcending the pure mechanics stipulated by the mathematical model of the game. Arguably, a real conductor, even if appropriately informed about the axioms of the game and its solutions, would still exhibit behaviours based not only on the utility function, but (presumably) also on aesthetic preference, feedback from the audience, the orchestra(s), the acoustic environment, and so forth. Notwithstanding the limitations of the mathematical framework that this article investigated, and given the new potential provided by machine learning and networked performance, game theory-based pieces continue to be exciting vehicles for structuring musical interaction and design that can open up to novel and unforeseen modalities of musical expression.

References

[1] Bang, Molly (1991). *Picture This:* Perception & Composition. Boston: Bulfinch Press.

[2] Bayer, Francis (1981). *De Schoenberg à Cage*. Paris: Kliencksieck.

[3] Beguš, Jelena Janković (2016). Playing the game with aleatorics and narrativity: *Linaia-Agon* by Iannis Xenakis. *New Sound*, v. 48, n. 2, pp. 109–130.

[4] Brackett, John (2010). Some Notes on John Zorn's Cobra. *American Music*, v. 28, n. 1, pp. 44–75.

[5] Brown, Tom; *et al* (2020). Language models are few-shot learners. In: Larochelle, H. *et al* (Ed.), *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901.

[6] Caillois, Roger (1961). *Man, Play and Games*. New York: Free Press of Glencoe.

[7] Cox, Christoph (2004). The Game Pieces. In: *Audio Culture:* Readings in Modern Music. New York: Continuum.

[8] DeLio, Thomas (1987). Structure and strategy: Iannis Xenakis' Linaia-Agon. *Interface*, v. 16, n. 3, pp. 143–164.

[9] Havryliv, Mark; Vergara-Richards, Emiliano (2006). From Battle Metris to Symbiotic Symphony: A New Model For Musical Games. In: *Proceedings of the 2006 International Conference on Game Research and Development*, Perth, Australia, pp. 260–268.

[10] Huizinga, Johan (1955). *Homo Ludens.* A study of the play-element in culture. Boston: Beacon Press.

[11] Kagel, Mauricio (1964). Match. Universal Edition.

[12] Kalonaris, Stefano (2016). Markov Networks for Free Improvisers. In: *Proceedings of the* 42*nd* *International Computer Music Conference*, Utrecht, The Netherlands, pp. 181–185.

[13] Kalonaris, Stefano (2017). Adaptive specialisation and music games on networks. In: *Proceedings of the* 13*th* *International Symposium on Computer Music Multidisciplinary Research*, Matosinhos, Portugal, pp. 420–429.

[14] Kalonaris, Stefano (2018a). Beyond Schemata in Collective Improvisation: A Support Tool for Music Interactions. *Leonardo Music Journal*, v. 28, n. 1, pp. 34–37.

[15] Kalonaris, Stefano (2018b). Satisficing goals and methods in human-machine music improvisations: Experiments with *Dory*. *Journal of Creative Music Systems*, v. 2, n. 2.

[16] Klein, Gary A.; *et al* (1993). *Decision Making in Action:* Models and Method. New York: Ablex.

[17] Lemke, Carlton E.; Howson, Jr. Joseph T. (1964). Equilibrium points of bimatrix games. *Journal of The Society for Industrial and Applied Mathematics*, v. 12, n. 2, pp. 413–423.

[18] Liuni, Marco; Morelli, Davide (2006). Playing Music: An installation based on Xenakis' musical games. In: *Proceedings of the Working Conference on Advanced Visual Interfaces*, AVI '06, New York, NY, USA, pp. 322–325.

[19] Morgenstern, Oskar; von Neumann, John (1947). *The Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

[20] Nowé, Ann; Vrancx, Peter; De Hauwere, Yann-Michaël (2012). Game theory and multi-agent reinforcement learning. In: M. Wiering & M. van Otterlo (Eds.), *Reinforcement Learning:* State-of-the-Art pp. 441–470. Berlin, Heidelberg: Springer Berlin Heidelberg.

[21] Osborne, Martin J.; Rubinstein, Ariel (1994). *A Course in Game Theory*. Cambridge: MIT Press.

[22] Shadow-Sky, Mathius (1980). Ludus Musicae Temporarium. http://centrebombe.org/livre/1980.b.html. Accessed: December 22nd, 2022.

[23] Sluchin, Benny (2005). Linaia-Agon: Towards an Interpretation Based on the Theory. In: *Proceedings of the International Symposium Iannis Xenakis*, pp. 299–311.

[24] Sluchin, Benny (2015). Linaia-Agon By Iannis Xenakis: Surpassing One's Limit / Le Dépassement De Soi. https://moderecords.com/catalog/284-xenakis/. Accessed: February 14th, 2022.

[25] Sluchin, Benny; Malt, Mikhail (2011a). Open form and two combinatorial musical models: The cases of Domaines and Duel. In: C. Agon *et al* (Ed.), *Mathematics and Computation in Music* pp. 255–269. Berlin, Heidelberg: Springer Berlin Heidelberg.

[26] Sluchin, Benny; Malt, Mikhail (2011b). Play and game in Duel and Strategy. In: *Proceedings of the Xenakis International Symposium*.

[27] Takagi, Hideyuki (2001). Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE*, v. 89, n. 9, pp. 1275–1296.

[28] Varga, Bálint András (2003). *Conversations with Iannis Xenakis*. London: Faber & Faber.

[29] von Stengel, Bernhard (2007). Equilibrium Computation for Two-Player Games in Strategic and Extensive Form. In: N. Nisan *et al* (Ed.), *Algorithmic game theory* pp. 53–78. Cambridge: Cambridge University Press.

[30] Watkins, Christopher J. C. H. (1989). *Learning from Delayed Rewards*. Ph.D. Thesis, Kings College.

[31] Xenakis, Iannis (1972). Duel. Éditions Salabert.

[32] Xenakis, Iannis (1992). *Formalized music:* thought and mathematics in composition. Hillsdale: Pendragon Press.